Noname manuscript No.
(will be inserted by the editor)

A Parallel Solution to Finding Nodal Neighbors in Generic Meshes

Gang Mei · Nengxiong Xu · Hong Tian · Shengwei Li

arXiv:1604.04689v3 [cs.CG] 5 May 2016

Received: date / Accepted: date

Abstract In this paper we specifically present a parallel solution to finding the one-ring neighboring nodes and elements for each vertex in generic meshes. The finding of nodal neighbors is computationally straightforward but expensive for large meshes. To improve the efficiency, the parallelism is adopted by utilizing the modern Graphics Processing Unit (GPU). The presented parallel solution is heavily dependent on the parallel sorting, scan, and reduction, and can be applied to determine both the neighboring nodes and elements. To evaluate the performance, the parallel solution is compared to the corresponding serial solution. Experimental results show that: our parallel solution can achieve the speedups of approximately 55 and 90 over the corresponding serial solution for finding neighboring nodes and elements, respectively. Our parallel solution is efficient and easy to implement, but requires the allocation of large device memory.

Keywords Computational Geometry · Mesh Topology · Neighbors Finding · Parallel Programming · GPU

1 Introduction

Mesh generation plays an important role in geometric modeling, computer graphics, and numerical simulations. After generating various types of meshes, typically mesh editing is intentionally performed to modify or improve the generated meshes to meet desired requirements. In mesh editing such as Boolean operations [1] or mesh optimization [2], the local mesh connectivity especially the adjacent/neighboring nodes and elements for each node or element is frequently needed to reduce the computational cost of local search.

The finding of one-ring nodal neighbors in arbitrary valid mesh is computationally straightforward, and can be completely carried out based on the connectivity/topology of meshes. The simplest method is to loop over all elements in a mesh to identify: (1) which pair of nodes is connected by an edge and (2) which nodes are contained in an element [2,3,4]. This is because that: (1) any pair of nodes connected by an edge is the one-ring neighboring node for each other; and (2) any element is directly the one-ring neighboring element for those nodes it contains.

Another simple method for finding the one-ring neighboring nodes for each vertex in a polygonal mesh was introduced by Dahal and Newman [5]. They first determined the boundary vertices by finding the opposite edges for each vertex and then forming a closed polygon using those opposite edges. They adopted the vertices of the closed polygon (i.e., the opposite edges) as the one-ring neighbors for each vertex.

Both of the above neighbors-finding methods are quite easy to implement in the serial programming pattern. However, due to the fact that it needs to loop over all the elements of a mesh in sequential, the computational cost is in general too high for large size of meshes; and this will reduce the computational efficiency of the entire mesh editing procedure.

An effective strategy to improve the efficiency of the neighbors-finding procedure is to parallelize it on various parallel computing architectures such as multi-core CPUs or

G. Mei · N. Xu · S. Li
Department of Geological Engineering, Qinghai University, Qinghai, China
School of Engineering and Technolgy,
China University of Geosciences, Beijing, China
E-mail: {gang.mei; xunengxiong; lishengwei}@cugb.edu.cn

H. Tian
Faculty of Engineering,
China University of Geosciences, Wuhan, China
E-mail: htian@cug.edu.cn

many-core GPUs, i.e., to design parallel solution to finding neighbors.

However, when finding nodal neighbors in parallel, there exists the race condition. The race condition issue appears when two different parallel threads may need to be written in the same memory position [6]. When looping over all the elements in a mesh to find the nodal neighbors, two neighboring nodes of the same vertex may be found concurrently within two parallel threads; and the indices of the two neighboring nodes may need to be written in the same memory position for storing. In this case, race condition arises.

Currently, the most commonly used method to address the above problem is to color the mesh first and then looping over those elements with the same color simultaneously to find neighbors [4,6,7,8,9]. This coloring-based method is very effective and efficient for large-scale meshes, and is quite suitable to be applied in parallel pattern. The only minor shortcoming is that: it is needed to color the mesh into several groups of elements and thus needs additional computational cost.

In this paper, without the use of mesh-coloring, we specifically design and develop a parallel solution to finding nodal neighbors by utilizing the power of modern GPUs. Our solution is efficient, simple and easy to implement, which heavily depended on the use of parallel primitives such as sorting, scan, and reduction. In addition, in our solution there is no need to adopt complex mesh data structures; and only arrays of integers are required. To evaluate the performance of our parallel solution, we compare it to the corresponding serial solution in several experiments.

2 Methods

2.1 Basic Ideas behind Our Solution

2.1.1 The Idea for Finding Neighboring Nodes for Each Vertex

In any valid meshes, any pair of neighboring nodes is connected using an edge. An edge can be represented with two nodes (and further, the indices of two nodes). A mesh typically has plenty of edges (i.e., pairs of nodes). When gathering all edges of a mesh (see Figure 1(a) and 1(b)), the edges can be stored in an array consisting of n pairs of integer values; see Figure 1(c). The array of edges can be also considered as two arrays of integers; see Figure 1(d). If adopting the first array of integers as the keys for sorting and the second array and the correspondingly attached values, the indices of all the neighboring nodes for the same vertex can be easily found by performing a parallel sorting; see Figure 1(e).

2.1.2 The Idea for Finding Neighboring Elements for Each Vertex

An element in a mesh is composed of several nodes (Figure 2(a)). Each element is by nature the one-ring neighboring element of those nodes it contains. A pair of integer values can be used to simply demonstrate this relationship: the first integer is the index of one of the nodes in an element; and the second is the index of the element itself; see Figure 2(b).

For an element, several such pairs of integers can be formed. And for an arbitrary mesh, a group of such pairs of integers can be obtained, and stored in two arrays of integers. Similar to the finding of neighboring nodes for each vertex, if adopting the first array of integers as the keys for sorting and the second array of integers as the correspondingly attached values, the indices of all the neighboring elements for the same vertex can be easily found by performing a parallel sorting; see Figure 2.

2.2 Implementation Details

Our solution is applicable to arbitrary meshes. However, to demonstrate our solution, here we specifically present the implementation details of our solution when applied to the triangular surface mesh.

Our solution is implemented by strongly utilizing the parallel primitives provide by the library *Thrust* [10] such as the parallel sorting and reduction. Thrust is a parallel algorithms library containing the efficient parallel primitives such parallel scan, reduction, and sorting. In our solution, the quite efficient primitive, parallel sorting is adopted to extremely fast sort the arrays of integers to find the neighboring nodes and elements. More implementation details are listed as follows.

2.2.1 The Finding of Neighboring Nodes

The process of finding the neighboring nodes for each vertex is as follows.

First, we create pairs of integers according to the edges in each triangle. Noticeably, we can form three edges (pairs of integers) for a triangle when the three nodes in the triangle are organized in count-clockwise (CCW) order and another three pairs when nodes are organized in clockwise (CW); see Figure 1(b). That is, a triangle can produce six pairs of integers. These edges/pairs can be obviously created in parallel. We specifically design a CUDA kernel to realize this; see the lines $7 \sim 8$ in Figure 3.

After creating the pairs of integers that are stored in two arrays of integers, the second step is to sort those pairs according to the first array of integers (see line 11 in Figure 3). This procedure can be extremely fast performed by using the specific function `thrust:: sort_by_keys()`.

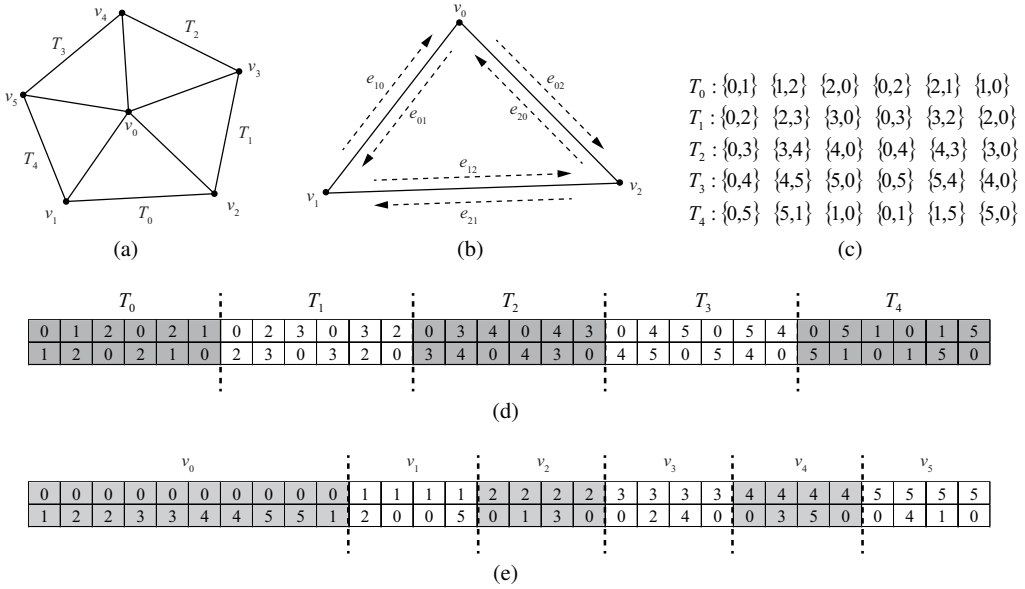

Fig. 1: A simple illustration of finding neighboring nodes for each vertex in a mesh

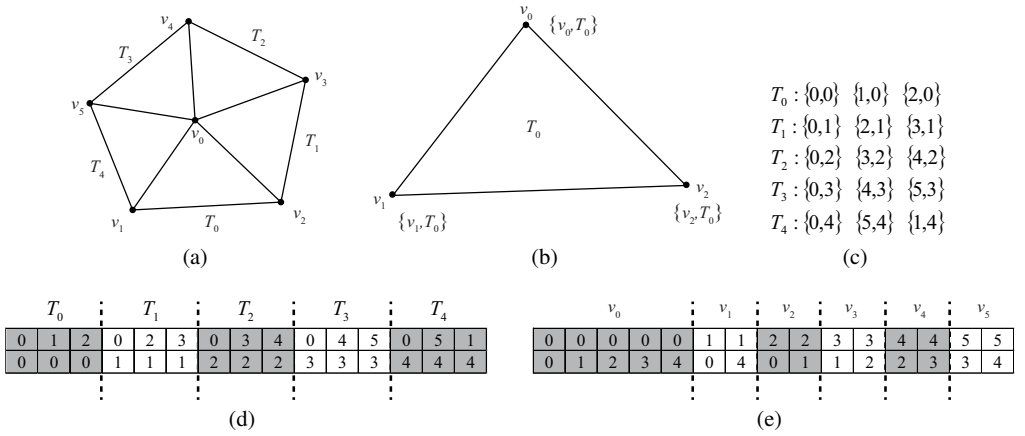

Fig. 2: A simple illustration of finding neighboring elements for each vertex in a mesh

The third step is to determine (1) the total number and (2) the detailed indices of the neighboring nodes for each vertex, which can be realized by using segmented scan and reduction. The ideas behind performing the segmented scan and reduction have been presented in our previous work [11].

To determine the number of neighbors, we first create a helper array containing the same value 1 (i.e., line 40 in Figure 3), and then perform a parallel segmented reduction by using the function `thrust::reduce_by_keys()`; see lines 42 $\sim$ 46 in Figure 3. To obtain the indices of the neighboring nodes, we also first create a helper array of sequenced integers (i.e., line 28 in Figure 3), and then perform a parallel segmented scan by using `thrust::unique_by_keys()`; see lines 27 $\sim$ 29 in Figure 3. After performing the segmented reduction and scan, both the number and indices of neighbors can be found and then transferred into other target arrays for further mesh editing such as Boolean operations or mesh optimization.

2.2.2 The Finding of Neighboring Elements

The process of finding the neighboring elements is quite similar to that of finding the neighboring nodes. The first step is also to form the pairs of integers (i.e., two arrays of integers), then to sort according to the first array of integers, and third use the parallel segmented reduction and scan to further determine both the total number and the detailed indices of the neighboring elements.

However, there is a remarkable difference between the process of finding neighboring nodes and elements. When

finding the neighboring elements, the first integer value of any pair is the index of a node in an element; and the second value of the pair of integers is the index of the element itself. In contrast, in the finding of neighboring nodes, both the first and the second integer value of any pair is the index of node.

3 Results and Discussion

In this section, we will evaluate the performance of our parallel solution by comparing to the corresponding serial solution when applied to the triangular surface meshes. Moreover, we will analyze the advantages and shortcomings of our parallel solution based on the experimental results.

3.1 Results

Six groups of experimental tests are carried out to evaluate the performance of our parallel solution. These experimental tests of the parallel solution are performed on the desktop computer which features with the NVIDIA GeForce GT640 (GDDR5) graphics card with 1GB memory and the GPU programming model CUDA v7.0. The experiments of the corresponding serial solution are performed on Windows 7 SP1 with an Intel i5-3470 CPU (3.2 GHz and 4 Cores) and 8GB of RAM memory.

These six triangular surface mesh models employed for testing are directly obtained from the Stanford 3D Scanning Repository (`http://www.graphics.stanford.edu/data/3Dscanrep/`) and the GIT Large Geometry Models Archive (`http://www.cc.gatech.edu/projects/large_models/`); see Figure 4 and Table 1. The execution time of both our parallel solution and the serial solution for finding the neighboring nodes and elements are listed in Table 1.

The experimental results listed in Table 1 indicate that: (1) our parallel solution is approximately 55 and 90 times faster than the serial solution when finding the neighboring nodes and elements, respectively; and (2) for the entire neighbors-finding procedure, our parallel solution achieves the speedup of approximately 60 over the serial solution.

3.2 Discussion

The finding of neighboring nodes and elements for each vertex in arbitrary meshes is computationally straightforward, but expensive for large meshes. In this paper, we specifically design and develop a parallel solution to improve the computational efficiency. In this section, we will analyze the advantages and shortcomings of our parallel solution.

3.2.1 The Advantages of Our Parallel Solution

The first advantage of our parallel solution is that: it is easy to implement due to the reason it does not need to perform the mesh-coloring before finding neighbors. The mesh-coloring technique is frequently used to deal with the race condition issue [6].

When finding nodal neighbors in parallel, there exists the race condition. The race condition issue appears when two different parallel threads may need to be written in the same memory position [6]. When looping over all the elements in a mesh to find the nodal neighbors, two neighboring nodes of the same vertex may be found concurrently within two parallel threads; and the indices of the two neighboring nodes may need to be written in the same memory position for storing. In this case, race condition arises.

Currently, the most commonly used method to address the above problem is to color the mesh first and then looping over those elements with the same color simultaneously to find neighbors [4,6,7,8,9]. This coloring-based method is very effective and efficient for large-scale of meshes, and is quite suitable to be applied in parallel pattern. The only minor shortcoming is that: it is needed to color the mesh into several groups of elements and thus needs additional computational cost.

In our parallel solution, we have redesigned the process of finding neighbors to avoid the use of mesh-coloring by strongly exploiting those efficient parallel primitives such as parallel sorting and scan. In addition, there are no complex data structures; and only arrays of integers are needed. Thus, our parallel solution is easy to implement in practice.

The second advantage of our parallel solution is the acceptable efficiency. The experimental results listed in Table 1 indicate that: our parallel solution can generate the speedups of approximately 55 and 90 over the serial solution when finding the neighboring nodes and elements, respectively.

This performance gains benefit from the parallelization performed on the GPU. By analyzing and reorganizing the process of finding nodal neighbors, we have transferred the entire process into several sub-procedures of parallel sorting, scan, and reduction, while these parallel primitives are extremely fast for the large size of data.

Another reason why our parallel solution is quite efficient is that: there are no complex data structures; and only arrays of integers are needed. Inherently, operations and computations for discrete arrays of integer values rather than arrays of structures are quite fast on the GPU. We specifically avoid using arrays of structures such as pairs, but chose to directly use the arrays of values. This leads additional performance gains.

```
// Find Neighboring Nodes in Parallel
// Step 1: Create edges / pairs of integers
thrust::device_vector<int> first_d (6 * nTrgl);   // 1st integer
thrust::device_vector<int> second_d(6 * nTrgl);   // 2nd integer
int * first_d_ptr   = thrust::raw_pointer_cast(&first_d[0]);
int * second_d_ptr  = thrust::raw_pointer_cast(&second_d[0]);
CreateEdges<<<blocksPerGrid_Trgl, threadsPerBlock>>>(
    first_d_ptr, second_d_ptr, nTrgl, trgls_d);

// Step 2: Sort by keys
thrust::sort_by_key(first_d.begin(), first_d.end(), second_d.begin());

// Step 3: Segmented parallel scan and reduction
// Prepare for parallel scan and reduction
thrust::device_vector<int> keys(6 * nTrgl);
thrust::device_vector<int> vals(6 * nTrgl);
thrust::fill(keys.begin(), keys.end(), 0);
thrust::fill(vals.begin(), vals.end(), 0);

thrust::device_vector<int> head_idx(nVert);
thrust::device_vector<int> num_neig(nVert);
int * head_idx_ptr = thrust::raw_pointer_cast(&head_idx[0]);
int * num_neig_ptr = thrust::raw_pointer_cast(&num_neig[0]);

// Find indices of neighbors using: unique by keys
// Need additional array of values {0, 1, 2, 3, ...}
keys = first_d;
thrust::sequence(vals.begin(), vals.end());
thrust::unique_by_key(keys.begin(), keys.end(), vals.begin());

int * keys_ptr = thrust::raw_pointer_cast(&keys[0]);
int * vals_ptr = thrust::raw_pointer_cast(&vals[0]);

Scatter_Copy_Kernel<<<blocksPerGrid_Vert, threadsPerBlock>>>(
    keys_ptr, vals_ptr, nVert, head_idx_ptr);  // Save

// Find numbers of neighbors using: reduce by keys
// Need additional array of values {1, 1, 1, 1, ...}
keys = first_d;
thrust::fill(vals.begin(), vals.end(), 1);

thrust::device_vector<int> tmp_keys(nVert);
thrust::device_vector<int> tmp_vals(nVert);

thrust::reduce_by_key(keys.begin(), keys.end(), vals.begin(),  // Input
                      tmp_keys.begin(), tmp_vals.begin());     // Output

thrust::copy_n(tmp_keys.begin(), nVert, keys.begin());
thrust::copy_n(tmp_vals.begin(), nVert, vals.begin());

Scatter_Copy_Kernel<<<blocksPerGrid_Vert, threadsPerBlock>>>(
    keys_ptr, vals_ptr, nVert, num_neig_ptr);  // Save
```

Fig. 3: Partial GPU implementation for finding neighboring nodes

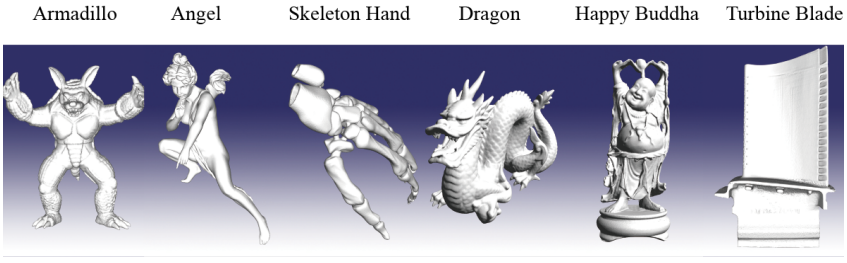

Fig. 4: Six mesh models that are employed for experimental tests

Table 1: Comparison of efficiency of our parallel solution and the serial solution (1k = 1000)

Mesh Model	Num. of Vertex	Num. of Triangle	Serial Time (/ms)		Parallel Time (/ms)		Speedup		
			Finding Node	Finding Elem.	Finding Node	Finding Elem.	Finding Node	Finding Elem.	Overall
Armadillo	172k	346k	2527	1732	50.0	22.1	50.5	78.4	59.1
Angel	237k	474k	3635	2580	66.7	29.4	54.5	87.8	64.7
Skeleton Hand	327k	655k	4806	3305	89.1	38.3	53.9	86.3	63.7
Dragon	437k	871k	6543	4488	112.7	48.5	58.1	92.5	68.4
Happy Buddha	543k	1088k	8577	6072	138.5	60.6	61.9	100.2	73.6
Turbine Blade	882k	1765k	13931	10000	221.3	93.3	63.0	107.2	76.1

3.2.2 The Shortcomings of Our Parallel Solution

Although our parallel solution is efficient and easy to implement, it has an obvious shortcoming, i.e., it requires more device memory than that of the serial solution. This additionally required device memory is allocated for performing the sorting, scan, and reduction.

In the serial solution, a STL (C++ Standard Template Library) container, `vector<int>`, is adopted to allocate an array to dynamically store the indices of neighboring nodes for each vertex. The size of the array can be dynamically determined without redundant space. Similar, another array of integers is needed to hold the indices of neighboring elements. Moreover, the number of neighboring nodes and elements are directly the size of the dynamic arrays, which can be easily and automatically determined. Thus, there is no need to allocate other additional arrays.

In contrast, in our parallel solution, six additional arrays of integers are required. First, two arrays of integers need to be allocated to store the edges (i.e., pairs of integers). Second, another two arrays of integers are needed to hold the first indices and numbers of neighboring nodes. And third, to perform the segmented parallel reduction and scan, another two temporary arrays are also required.

Due to the required additional arrays of integers, our parallel solution cannot be applied to quite large size of meshes since the device memory (the global memory) of most current GPUs is quite limited. Thus, future work may focus on redesigning the parallel process of finding nodal neighbors to reduce the device memory.

4 Conclusions and Outlook

In this paper, we have designed and developed a parallel solution to finding the neighboring nodes and elements for each vertex in an arbitrary mesh by exploiting the GPU. Our solution is a topology-based method, and is heavily dependent on the use of parallel sorting, scan, and reduction. We have compared our parallel solution to the corresponding serial solution to evaluate its performance. We have found that: our parallel solution is approximately 55 and 90 times faster than the corresponding serial solution when finding the neighboring nodes and elements, respectively. Our solution is efficient, simple and easy to implement, and can be applied for arbitrary meshes. However, our parallel solution requires the allocation of large device memory; and thus future work is planned to be carried out to address this problem. To benefit the community, the complete source code of our solution is publicly available for any potential usages.

Acknowledgements This research was supported by the Natural Science Foundation of China (Grant Numbers: 40602037, 40872183, and 51541405), China Postdoctoral Science Foundation (2015M571081), and the Fundamental Research Funds for China Central Universities (2652015065). The authors would like to thank the editor and the reviewers for their contributions on the paper.

References

1. F. Feito, C. Ogayar, R. Segura, M. Rivero, Computer-Aided Design **45**(3), 705 (2013). DOI 10.1016/j.cad.2012.11.004
2. J. DAmato, M. Vnere, Journal of Parallel and Distributed Computing **73**(8), 1127 (2013). DOI 10.1016/j.jpdc.2013.03.007
3. G. Mei, J.C. Tipper, N. Xu, Arabian Journal for Science and Engineering **39**(11), 7907 (2014). DOI 10.1007/s13369-014-1406-y
4. J. Chen, X. Jin, Z. Deng, The Visual Computer **31**(2), 119 (2014). DOI 10.1007/s00371-014-0924-7
5. S. Dahal, T.S. Newman, in *IEEE SOUTHEASTCON 2014* (2014), pp. 1–7. DOI 10.1109/SECON.2014.6950720
6. L.S. Duarte, W. Celes, A. Pereira, I.F. M. Menezes, G.H. Paulino, Structural and Multidisciplinary Optimization **52**(5), 845 (2015). DOI 10.1007/s00158-015-1252-x
7. R.Y.G.Z.L.O. Eric Shaffer, Zuofu Cheng, Advances in Parallel Computing **25**, 285 (2013). DOI 10.3233/978-1-61499-381-0-285
8. D. Benítez, E. Rodríguez, J.M. Escobar, R. Montenegro, *Proceedings of the 22nd International Meshing Roundtable* (Springer International Publishing, Cham, 2014), chap. Performance Evaluation of a Parallel Algorithm for Simultaneous Untangling and Smoothing of Tetrahedral Meshes, pp. 579–598. DOI 10.1007/978-3-319-02335-9_32
9. Z. Cheng, E. Shaffer, R. Yeh, G. Zagaris, L. Olson, Engineering with Computers pp. 1–10 (2015). DOI 10.1007/s00366-014-0393-7
10. N. Bell, J. Hoberock, in {*GPU*} *Computing Gems Jade Edition*, ed. by W.m.W. Hwu, Applications of GPU Computing Series (Morgan Kaufmann, Boston, 2012), pp. 359 – 371. DOI 10.1016/B978-0-12-385963-1.00026-5
11. G. Mei, N. Xu, L. Xu, arXiv:1601.05904 (2016). URL `http://arxiv.org/abs/1601.05904`